REPORT DOCUMENTATION PAGE		*Form Approved* *OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
29-08-2011	Final Report	From March 1, 2008 to May 31, 2011

4. TITLE AND SUBTITLE An Optimization Framework for Air Force Logistics Models	**5a. CONTRACT NUMBER**
	5b. GRANT NUMBER FA9550-08-1-0369
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Retsef Levi and Tom Magnanti	**5d. PROJECT NUMBER**
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MIT	**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Optimization and Discrete Mathematics Program in the Mathematics, Information and Life Sciences, Directorate of the Air Force Office of Scientific Research (AFOSR)	**10. SPONSOR/MONITOR'S ACRONYM(S)** AFOSR
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-OSR-VA-TR-2012-0873

12. DISTRIBUTION/AVAILABILITY STATEMENT

Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The goal of this project was to create a new framework of models and algorithms to address grand challenges in the US Air Force logistics. We have developed new models, algorithms and performance analysis techniques that address major current and future logistics issues within the Air Force. In particular, we have focused on strategic and operational issues in the management of maintenance resources and its interactions with inventory and transportation considerations. Over the last 50 years the Air Force maintenance resources has significantly decreased and become a critical bottleneck to the Air Force ability to accomplish its missions. However, relatively little attention has been given to maintenance considerations compared to other considerations, such as inventory management. The algorithms that we developed are provably near-optimal and can be applied to realistic large scale instances. Some of the work that has been done in this project has promising potential to impact and improve some of the current Air Force logistics best practices, and ultimately lead to decision support tools that could assist decision makers within the Air Force. We believe that this work creates new avenues for additional basic research related to Air Force logistics core current and future problems.

15. SUBJECT TERMS

Optimization, logistics, inventory, performance analysis, worst-case, transportation, online algorithms

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			**19b. TELEPHONE NUMBER** *(Include area code)*

FA9550-08-1-0369 - An Optimization Framework for Air Force Logistics Models
Final Report - August 29, 2011

1 Introduction

In this report we outline the work that has been performed as part of grant FA9550-08-1-0369 awarded to us by the *Optimization and Discrete Mathematics Program* in the Mathematics, Information and Life Sciences Directorate of the *Air Force Office of Scientific Research* (AFOSR). The report covers the time period March 1, 2008 through May 31, 2011. We shall summarize the body of work that has been performed, list the papers that have been published, explain how the funding has been used and also describe the potential impact on the Air Force.

As the name of the grant suggests, the goal of this project was to start creating a new framework of models and algorithms to address grand future challenges of the US Air Force logistics. To identify current and future challenges and guide the basic research done within the project, we have engaged in a broad range of working relationships with various central Air Force logistics units and personnel (see detailed list in the Appendix below). The contributions of the work done under this grant are two-fold. First, we have developed several new models, algorithms and analysis techniques that address major current and future logistics issues within the Air Force. We believe that this work creates new avenues for additional basic research related to Air Force logistics core current and future problems. Second, some of the work that has been done in this project has promising potential to impact and improve some of the current Air Force logistics best practices, and ultimately lead to decision support tools that could assist decision makers within the Air Force. In addition, the funding of this grant was partially used to promote more formal collaborations between academics and Air Force logistics decision makers.

Based on the discussions we had with several senior Air Force logistics personnel, we have identified several themes of issues that are central to important current and future logistics challenges faced by the US Air Force. This has guided the focus of the work in this project. Next we motivate and describe the main themes of work performed within this project:

Theme (I): Maintenance Management - Optimization Models and Algorithms. The management of maintenance resources has become a major challenge in the US Air Force logistics system and a major enabling capability for the Air Force ability to successfully accomplishing its various missions. There are several trends that have affected the increasing importance of maintenance management. First, many, if not all, the current and future Air Force weapon systems are based on modular engineering designs that allow the maintenance of some components in the system without affecting other components. This creates some major opportunities for cost reductions and more efficient use of maintenance resources, but at the same time requires more sophisticated maintenance scheduling and planning approaches to effectively exploit the modular structure. Second, over the last 5 decades there has been a significant decrease of over 60% in the number of Air Force logistics personnel (see Airman Magazine –The Book 2010), particularly, front-line logisticians. This decrease in the respective personnel levels has made maintenance resources a major bottleneck for the Air Force logistics operations. Another contributing factor to this latter trend was the fundamental change in the Air Force missions, which shifted from the cold war scenario that focused on one primary theater to a scenario, in which the Air Force has to simultaneously manage multiple geographically distant theaters (e.g., Iraq and Afghanistan). This shift has also made centralization of maintenance resources far more costly and challenging.

Traditionally, the focus of the Air Force as well as the academic work on Air Force logistics issues has been primarily on *inventory management* issues, particularly, the effort to ensure the availability of *spare parts*. In comparison, relatively very little attention was dedicated to the management of maintenance resources. While spare parts availability is still a central issue, maintenance aspects play a role that is equality important and critical to the Air Force mission capabilities. In fact, many of the parts being used in critical weapon systems, particularly aircrafts, are reusable in the sense that after being used for a predefined period of time, they have to go through maintenance, after which they can be used again. This creates very challenging tradeoffs between inventory, maintenance and transportation management. These tradeoffs are currently not well understood, particularly in the context of the future missions and weapon systems of the Air Force. There is also an important complicating organizational factor that makes these tradeoffs even more challenging. Specifically, the management of maintenance and inventory (spare parts) and the management of transportation resources are currently slitted between different organizations, the *Air Force Material Command* (AFMC) and the *Air Force Mobility Command* (AMC), respectively.

In this theme of work, we have developed several new models that capture major maintenance tradeoffs both at the initial design level and the operational level. The operational models capture some of the major interactions between maintenance management and inventory and transportation management related considerations, and provide decision makers a framework to think about all of these issues and tradeoffs in a much more integrative way. Moreover, we have developed provably near-optimal algorithms to solve the respective models in a way that could inform and guide decision makers in the Air Force logistics system. The algorithmic and performance analysis techniques that we have developed have a promising potential to be applicable in other logistics related models.

Theme (II): Online and Real Time Models and Algorithms. In many practical scenarios Air Force logistics management related decisions are made under many uncertainties. For example, there are many uncertainties around the future operational scenarios, equipment and weapon system failure rates and many other important parameters. Modeling uncertainty in these settings is particularly challenging, whereas two of the major issues are the manner by which the uncertainty is modeled and the computational complexity of the model. In many of the US Air Force related practical settings, decision are made under severe time constraints, specifically, they have to be made in real-time or within a very short period of time (say within minutes, seconds or even less). Moreover, the Air Force logistics system is very large scale and involves many different weapon systems, operational units, geographical locations and many other constraints. These aspects make the respective challenges even more complex. Specifically, future decision support tools for these settings must be based on *real-time and online algorithms* that can run very fast on large scale instances, but at the same time incorporate the potential impact of uncertainty. As mentioned above modeling uncertainty is a major challenge in creating effective decision support tools for the management of large scale logistics systems. For example, one way to model uncertainty is to assume that the respective probability distributions are given as part of the input. However, in many settings it is not realistic to expect to have a reliable specification of the underlying probability distributions. In these scenarios, one could take a *data-driven* approach and assumes that there is only historical data available to the decision maker. Unfortunately, in many settings there is very little historical data, and even the data available might not be representative of the future state of the system. To model situations like this, one could apply an *online* approach, in which there are minimal/no assumptions on the future evolution of the system. The way uncertainty is modeled has a direct impact on the computational complexity of the model. Specifically, if one wishes to account for all possible future scenarios the resulting algorithms could be very slow and potentially even not tractable; this is known as the *curse of dimensionality*. In addition, the way uncertainty is modeled has direct impact on the quality of the solution. On one hand, a misspecified model could lead to bad solutions, but on the other hand, very general way to model uncertainty could lead to very conservative solutions that do not exploit available information about the future.

In this theme of work, we have developed several new models and algorithms that capture logistic planning and management under uncertainty. Moreover, our modeling framework allow flexibility regarding the way uncertainty is modeled. All the algorithms that we have developed have both provably a-priori and posteriori performance guarantees. That is, we obtain *a-priori worst-case guarantees* that imply that our algorithms are guaranteed to perform near optimal, and in addition, for each specific instance the algorithms provide a certificate that indicate how close they perform compared to the best possible on the specific input instance.

2 Major Results and Contributions

In this section, we describe the major results and contributions obtained in this project.

Models for maintenance of modular systems. In joint work with with Jack Muckstadt, Danny Segev and Major Eric Zarybnisky (an Air Force PhD student co-supervised by the PIs), we studied various aspects in the management of *scheduled maintenance activities for modular systems*, such as an aircraft engine. These systems consist of components with *cycle limits* that specify the maximum number of periods of use between subsequent maintenance actions. For example, a cycle for the starter system in an aircraft engine could be one startup sequence. For components in an aircraft braking system, a cycle could be one landing sequence. Each component can be used for a certain number of cycles and then must be repaired or replaced due to safety or failure concerns. These cycle limits are determined through a number of methods including physical testing, simulation, and analytical assessment. Although it is possible that components fail prior to their cycle limits, due to the conservative nature of these cycle limits, such events are extremely rare. As a result, it is common to assume that a component is operational until its cycle limit is reached and that after maintenance it again has a full cycle limit. While much of the literature has examined stochastically failing systems, preventative maintenance of *usage limited* components has received less attention.

In the *modular maintenance scheduling problem* [8, 11] we have studied a single modular system that consists of components with associated cycle limits. The goal is to compute a feasible maintenance schedule that minimizes the cost associated with component maintenance. This model captures some fundamental tradeoffs in the design phase of the system. By making cost tradeoffs early in development, program managers, designers, engineers, and test conductors can better balance the up front costs associated with system design and testing with the long term cost of maintenance. The model provides a framework for design teams to evaluate different design and operations concepts and then evaluate the long term costs. (For example, how to tradeoff the additional cost of extending the cycle-limit of a given component in the system during the development phase versus the long term maintenance cost savings.)

The typical cost structures that arise in practical settings are submodular in the subset of components being maintained. However, finding the optimal policy under these assumptions is computationally challenging. Moreover, the optimal policy can be very complex and does not provide the operational simplicity that is essential for implementation in many practical settings. A natural approach that is often considered in practice is to use cyclic policies that maintain each component at a fixed frequency. The question that arises is the increase in cost associated with using potentially suboptimal cyclic policies. We develop two algorithms to compute provably near-optimal cyclic policies. The *cycle rounding* algorithm computes the cycle-limits iteratively by considering the components in the system in increasing cycle-limits. The algorithm provides a worst-case approximation guarantee of 2. That is, for any input instance of the problem, the algorithm computes a cyclic policy with cost that is at most twice the optimal (over all policies, not necessarily cyclic policies). We also develop a class of *shifted power-of-two* algorithms. The cycle-limits are rounded to power-of-two times a shifting parameter and then an optimal policy with respect to the rounded cycle-limits is computed. (It can be shown that if all cycle-limits are power-of-two's the optimal policy is to maintain each component exactly when it is due.) Based on an innovative cost decomposition scheme and randomized analysis, we show that there is a small set of shifted power-of-two policies, the best of which is guaranteed to have cost at most $1/\ln(2)$ times the optimal cost. This guarantee holds for any submodular increasing cost function. Interestingly, the set of shifted power-of-two policies can be computed based only on the cycle limits independent of any other parameter of the problem, including the cost functions. Moreover, we show that one can choose an a-priori set of shifted power-of-two policies entirely independent of *any* parameter of the problem, and obtain constant worst-case guarantees. The guarantees are improving as the size of the this set grows larger and converge (quickly) to $1/\ln(2)$. In fact, even with just a few policies that are chosen a-priori, on can obtain a provable worst-case guarantee very close to $1/\ln(2)$. The analysis is obtained based on innovative linear programming approach that, for any given predefined number of chosen policies, reveals the worst-case guarantee, as well as tight worst-case instances. This is quite surprising in light of the fact that the policies are chosen with no knowledge on the specific input instance, but provide the guarantees for any instance with general submodualr cost functions. In extensive computational experiments, these cyclic policies perform extremely within a few percentages of optimal and much better than the worst-case guarantees.

Once a modular system has moved into operations, manpower and transportation scheduling become important considerations when developing maintenance schedules. To address the operations phase, we develop the *modular maintenance and system assembly model* [10, 9] to balance the tradeoffs between inventory, maintenance capacity, and transportation resources. This model explicitly captures the risk-pooling effects of a central repair facility while also modeling the interaction between repair actions at such a facility. The full model is intractable for all but the smallest instances. Accordingly, we decompose the problem into two parts, the *system assembly* portion and *module repair* portion. Even the decomposed models are in general computationally challenging. We first study several practically interesting special cases, and develop optimal algorithms and heuristics to solve them. Finally, we use the output of these models together with the algorithms developed for the modular maintenance scheduling problem to propose an integrated methodology for design and operations of these complex systems.

Maintenance of low observable aircrafts. In joint work with 1st Lt. Phil Cho (Master student co-supervised by the PIs), Vivek Farias and Major Eric Zarybnisky [2], we have studied the maintenance and flight scheduling of *low observable (LO) aircrafts*. The newest generation of fighter aircrafts in the Air Force, such as the F-22, has low-observable (LO) technologies that make them invisible to radar. This presents unique maintenance issues that did not exist for previous generations of fighter aircraft. In particular, the outer surfaces of the LO aircraft are coated with a metallic paint that is designed to minimize the radar signature of the aircraft. While LO aircraft have many design features that contribute to the LO capability of the aircraft, such as the shape and angles of the aircraft, the special outer metallic coating is the primary contributor to the increased maintenance requirements for LO aircraft. If the coating is damaged in any way, the radar signature of the aircraft can be affected. Since LO aircraft are not considered to be fully mission capable (FMC) unless their radar signature is below a certain level, maintenance personnel must continuously repair the metallic coating on LO aircraft in order to sustain an acceptable FMC rate for a fleet of aircraft.

An aircrafts LO capabilities degrade randomly, both in amount and physical location over the aircraft's body due to flying activities. Simple scrapes and dings can have a significant impact on the overall radar signature of an aircraft. Depending on the size, location, and shape of each specific damage, the overall impact of a single damage can range from being negligible to causing the aircraft to no longer be FMC. Each time an aircraft flies, maintenance personnel record all new damages into the *signature assessment system* (SAS). Particularly, each aircraft is associated with an evolving SAS number that reflects the estimated cumulative impact of all the damages on its LO capability. (Higher SAS number implies that the radar signature of the aircraft in higher.) Moreover, a SAS number higher than a certain threshold implies that the aircraft is not FMC. Therefore, LO maintenance personnel must carefully track the damages on each aircraft and decide when to repair them; this type of repair is . This type of maintenance is called *redux LO maintenance*. The effectiveness of a maintenance action depends on its length (i.e., for how many days the aircraft is put into maintenance), as well as the physical distributions of the damage to the aircraft coating. For example, if there is one major hit in a single location (sometimes called *'heavy hitter'*), maintenance will obtain a higher reduction in the SAS number compared with an aircraft with the same SAS number, but with multiple small hits over different locations. LO redux maintenance resources are scarce, and logistics personnel face major operational challenges in deciding which aircrafts should be put into maintenance to keep the entire feet at high level of FMC over time.

Currently, the decision process regarding SAS redux is largely dependent on the personalities of each maintenance unit. Since there is little published guidance regarding LO maintenance, each maintenance unit has the flexibility to make LO maintenance decisions however they see fit. Therefore, the various LO maintenance policies used by flying units throughout the Air Force can vary. In speaking with several experienced maintenance personnel, the LO maintenance decision process was described as being somewhat haphazard.

In this work, we model the LO SAS maintenance scheduling problem based on real data that records the SAS number evolution over 2.5 years (overall 5000 data points). The decision making problem is modeled as a variant of the *restless multi-armed bandit* problem. (This is a well-known model in stochastic control.) In addition, we use index policies that allow maintenance schedulers to quickly rank the aircrafts in the fleet based on each aircrafts current LO capability state, and use the ranking to decide which aircraft to enter into LO maintenance and for how long. We employ two algorithms to compute good index policies. In addition to maintenance scheduling, we explore policies for choosing which aircraft to fly to meet sortie requirements with a focus on the LO implications of those decisions. Finally, in extensive computational experiments we have demonstrated the strength of the index policies and the importance of the flight decision. In particular, we show that the index policies paired with good flight decisions perform within 10% of a computational upper bound. Several Air Force personnel are interested in testing the policies that we propose within a fidelity simulation setting they are developing to model various strategic logistics systems of the Air Force.

Supply chain management and logistics models with online customer selection. In joint work with Adam El-Machtoub the PI Levi consider new *online* versions of supply chain management and logistics models, where in addition to production decisions, one also has to make decisions regarding which customers (missions) to serve [4, 3]. Specifically, customers (missions) arrive sequentially over time during a *planning phase*, and the decision maker has to permanently decide whether to accept or reject each customer upon arrival. If rejected upon arrival, a lost-sales (or penalty) cost is incurred. Once the selection decisions are all made and the planning phase is over, one has to satisfy (serve) all the accepted customers (missions) with minimum possible production cost. The goal is to minimize the total lost-sales (penalty) and production costs. In contrast to previous work, our assumption is that customers arrive in an online manner. That is, upon arrival of a customer, the decision maker has information on all the customers that arrived prior to the current customer, but very limited or no information about future arrivals. In particular, we assume that the sequence of customer arrivals is generated by a *worst-case adversary*. (This is a very common assumption in the literature on online algorithms that have been applied to many combinatorial optimization models.)

We developed several novel algorithms for the respective online models that capture among others many variants of core logistics and supply chain optimization problems, such as the *single lot-sizing problem*, the *joint replenishment problem*, the *facility location problem* and *network design problems*. Our algorithms are based on repeatedly solving offline sub-problems. In particularly, upon each customer arrival, we solve a certain offline problem with respect to all of the customers arrived thus far, ignoring future arrivals and previously made decisions by the online algorithm. The solution of this offline problem informs the algorithm whether to accept the current customer. We then analyze the *competitive ratio* of the algorithm. That is, the performance of the online algorithm is compared to the performance that could be obtained in the entire sequence of customer arrival has been known upfront. In fact, the assumption is that the worst-case adversary will generate a sequence of demands to make the ratio between the cost of the solution produced by the online algorithm and the one of the optimal offline solution as high as possible. This is a very stringent bench mark that is very common in the analysis of online algorithms. However, we show that the algorithm that we propose have optimal or close to optimal competitive ratio guarantees. In computational experiments the algorithms perform very close to optimal, significantly better than the worst-case guarantees.

The models and the algorithmic and analysis techniques that were developed in this work has a promising potential to be applied in other settings, both to logistics models, as well as other application domains.

Other work. In addition to the above mentioned results and contributions the respective grant partially funded several other research efforts. In joint work with Con Shi (PhD student, partially funded by the grant) the PI Levi developed a new randomized cost-balancing policy for the *stochastic lot-sizing problem* with general demand structures. This is the first randomized policy applied to inventory management problems, and also the first policy to this core problem that has a worst-case performance guarantee [14].

In joint work with Joseph Geunes, Edwin Romeijn and David Shmoys, the PI Levi developed new logistics, inventory and location models with (offline) market selection, in which demands can be ignored (rejected) at the cost of a corresponding penalty. We developed a new framework that leveraged constant approximation algorithms for the classical variants of these models to these more general models [5].

In joint work with Niv Buchbinder, Tracy Kimbrel, Konstantin Makarychev and Maxim Sviridenko the

PI Levi developed several new online algorithms for core logistics models that have analytical competitive ratio [1]. The algorithms are based on a novel *primal-dual* approach that is based on a linear programming relaxation of the problems and use the dual of the linear program to guide the online algorithm.

In joint work with Tim Huh, Paat Rusmevichientong and Jim Orlin [7] and with Georgia Perakis and Joline Uichanco [12], the PI Levi studied *data-driven* variants of core stochastic inventory management models, in which unlike the traditional assumption that the demand distribution is given as part of the input, we assume that only historical demand or sales data is available. We develop several data-driven algorithms, and show analytically that they perform close to the optimal policy that could be computed if the demand distributions are known.

For other publications partially funded by this grant please see the reference list [6, 13, 15, 17, 16].

3 Other Activities

As part of the work funded by this award, we have developed a broad range of working relationships with key units and decision makers within the Air Force logistics system; a summary of the respective contacts is given in the Appendix. In addition, on October 2010, we organized a workshop at MIT on "Air Force Future Logistics Challenges: Decision-Support Models and Tools", to which we invited about 40 leading academics in the Operations Research community that do logistics related research work, as well as key decision makers from the Air Force logistics units. The workshop has been very successful and created a lot of promising engagements between academics and Air Force logistics decision makers.

4 Funding

Most of the funding provided by this award was allocated to support graduate students. The grant fully supported two Air Force students (1 PhD and 1 Master), and additional student was partially supported. In addition, the award partially supported a postdoctoral doctoral fellow (50% effort for one year). The PIs Levi and Magnanti used the grant to pay some of their summer support. The grant was also used for travel expenses and partial summer support for professor Jack Muckstadt from the School of Operations Research and Information Engineering in Cornell University. Professor Muckstadt is a retired Air Force logistic officer with over 40 years of experience with academic and practical work related to Air Force logistics. In addition to funding the research team the award was used to fund the workshop mentioned above.

A List of Air Force Contacts

1. HAF:
   - Mr. John W McDonald, AF/A4ID, John.McDonald@pentagon.af.mil
   - Mr. Eric R Nelson, AF/A4IT, Eric.Nelson@WPAFB.AF.MIL
2. AFLMA:
   - Capt. Christopher Arendt, Chief, Operational Analysis Branch, Air Force Logistics Management Agency, Comm: 334-416-4126 DSN: 596-4126, Christopher.Arendt@maxwell.af.mil
   - Dr. Doug Blazer, LMI, DSN 596-1406 Comm 334 416 1406, doug.blazer.ctr@maxwell.af.mil
   - Maj. Christopher A. Boone, Chief, Wargame Division, Air Force Logistics Management Agency, Maxwell-Gunter Annex, Maxwell AFB, DSN: 596-2325, Comm: 334-416-2325
   - 2nd Lt Eric Almeida, AFLMA/LGY, eric.almeida@maxwell.af.mil
   - Capt. Brian D Waller, AFLMA/LGLM, Chief, Maintenance Branch, Air Force Logistics Management Agency, DSN 596-4593/Comm 334-416-4593, Brian.Waller@maxwell.af.mil
   - Capt. Craig A. Lane, Chief, Wargaming Branch, Air Force Logistics Management Agency (AFLMA), 501 Ward Street, Maxwell AFB (Gunter Annex), AL 36114, DSN: 596.5803 Comm: 334.416.5803, craig.lane@maxwell.af.mil
3. AFMC:
   - Mr. Bob McCormick (AFMC/A9), HQ AFMC/A9A, Comm: 937-257-6920 (DSN: 787-6920), FAX: 937-656-1498 (DSN: 986-1498), Bob.McCormick@wpafb.af.mil
   - Mr. Mike McClure (AFGLSC), Operations Research Analyst, AFGLSC 402 SCMS/GUSB, Supply Chain Diagnostics and Analysis Flight, Bldg 70, Post 31, 5215 Thurlow Dr, ste. 5, WPAFB, OH 45433, DSN 986-2501 / Comm 937-656-2501, gordon.mcclure@wpafb.af.mil
   - Dr. Siva S. Banda, Director, Control Science Center of Excellence, Air Force Research Laboratory, AFRL/RBCA, 2210 Eighth St., Bldg. 146, Wright-Patterson AFB, OH 45433-7531, E-mail: siva.banda@wpafb.af.mil, Phone: (937)255-8677
   - Mr. Tim Meixner, ECSS Plan IPT Lead, 937-238-2741 (cell), 937-904-0795 (office)
   - Mr. Richard A. Moore, Chief, Studies and Analyses Division, HQ AFMC/A9A, DSN: 787-4044, COMM: 937-257-4044, Richard.Moore@wpafb.af.mil
   - Mr. Mike Niklas, Operations Research Analyst, Mike.Niklas@wpafb.af.mil
4. AMC:
   - Mr. Donald R. Anderson (AMC/A9), YD-03, DAF, Assistant Director, Analyses, Assessments, and Lessons Learned (AMC/A9), DSN: 770-7629 Comm: (618) 220-7629, Donald.Anderson-02@scott.af.mil
   - Dr. Donald R. Erbschloe, Chief Scientist, HQ AMC/ST, 402 Scott Dr., Unit 3EC, Scott AFB IL 62225, Voice: 618-229-4825 (DSN 779), Fax: 618-256-2502 (DSN 576), don.erbschloe@scott.af.mil
5. ACC:
   - Capt. John V. Miller, A9 Executive Officer, ACC DSN 574-5751, (757) 764-5751, 61X FAM, Air Combat Command, UFPM, ACC A9, John.Miller@langley.af.mil
   - Capt. Eric Gons, Operations Research Analyst, HQ ACC/A9A, 574-7684, (757) 764-7684, Eric.Gons@langley.af.mil
   - Capt. Elim H. Sniady, VaANG, Maintenance Operations Officer, 192d Maintenance Squadron, DSN: 574-0310, COMM: (757) 764-0310, CELL: 757-813-1903, elim.sniady@langley.af.mil

- 2nd Lt Alex J. Shamp, VaANG, OIC, LO CRF, 192d MXS, Langley AFB, VA 23665, DSN 575-5700, COMM (757) 225-5700, CELL (757) 869-0805, alex.shamp@langley.af.mil
- Maj. Lucius A. Cattles, 1 MOS, MX Operations Officer, Langley AFB VA, DSN: 574-6238 COMM: 757-764-6238, Lucius.Cattles@Langley.AF.Mil

6. PM-JAIT:

- Jerry Rodgers, PM J-AIT Support Contractor, 2461 Eisenhower Ave., Hoffman Building 1, Room 1140/1148, Alexandria, VA 22331-3011, (703) 325-2988, jerry.d.rodgers@us.army.mil, www.ait.army.mil
- Sumra Manning, CISSP, PMP, Engineering Manager, RFID, Federal Systems, Unisys, 11109 Sunset Hills Road, Reston VA 20190, Office 703 439 6580, Cell 215-527-8914, sumra.manning@unisys.com, sumra.manning@us.army.mil
- Lee James III, PM J-AIT, 703.373.1355, lee.jamesiii@us.army.mil

References

[1] N. Buchbinder, T. Kimbrel, R. Levi, K. Makarychev, and M. Sviridenko. Online make-to-order joint replenishment model: Primal-dual competitive algorithms. Working paper. Preliminary extended abstract appeared in SODA 2008, 2011.

[2] P. Cho, V. Farias, R. Levi, T. Magnanti, and E. Zarybnisky. Maintenance and flight scheduling of low observable aircraft. Working paper, 2011.

[3] A. Elmachtoub and R. levi. Online algorithms based on cost sharing schemes. Working paper, 2011.

[4] A. Elmachtoub and R. levi. Supply chain management and logistics models with online customer selection. Working paper, 2011.

[5] J. Geunes, R. Levi, E. Romeijn, and D. B. Shmoys. Inventory and facility-location models with market choice. Forthcoming in Mathematical Programming, 2011.

[6] V. Goyal, R. Levi, and D. Segev. Near-optimal algorithms for the assortment planning under dynamic substitution and stochastic demand. Submitted to Operations Research, second revision requested. Extended abstract appeared in MSOM 2009, 2010.

[7] T. Huh, R. Levi, P. Rusmevichientong, and J. Orlin. Adaptive data-driven inventory control with censored demand based on Kaplan-Meier estimator. Forthcoming in Operations Research, 2011.

[8] R. Levi, T. Magnanti, J. Muckstadt, D. Segev, and E. Zarybnisky. Maintenance scheduling for modular systems models and algorithms. Submitted to Operations Research, 2011.

[9] R. Levi, T. Magnanti, J. Muckstadt, D. Segev, and E. Zarybnisky. Multi-echelon engine and lru maintenance. Working paper, 2011.

[10] R. Levi, T. Magnanti, J. Muckstadt, D. Segev, and E. Zarybnisky. Multi-echelon modular maintenance and system assembly with transportation and maintenance capacities. Working paper, 2011.

[11] R. Levi, T. Magnanti, D. Segev, and E. Zarybnisky. The graph visiting problem. Soon to be submitted to Mathematical Programming, 2011.

[12] R. Levi, G. Perakis, and J. Uichanco. The data-driven newsvendor problem: New bounds and insights. Submitted to Operations Research, second revision was requested, 2011.

[13] R. Levi and A. Radovanovic. Provably near-optimal lp-based policies for revenue management in systems with reusable resources. *Operations Research*, 58(2):503–507, 2010.

[14] R. Levi and C. Shi. Approximation algorithms for stochastic lot-sizing inventory control models. In second revision in Operations Research. Extended abstract appeared in MSOM 2009, 2011.

[15] R. Levi and C. Shi. Revenue management of reusable resources with advanced reservations. Submitted to Operations Research, 2011.

[16] R. Levi, M. Sviridenko, and L. Yedidson. Lp-Based approximation algorithms for the multiitem lot-sizing problem with nonuniform capacity constraints. Working paper, 2011.

[17] R. Levi and L. Yedidson. The assebly inventory probelm is NP-Hard. Working paper, 2011.